Hotplug Device Does Not Show Correctly in Guest OS

Table of contents

Problem

After creating a hotplug device from the BlueField side, probing virtio drivers does not create the virtio-net device correctly.

Solution

The problem may happen due to the following reasons.

BAR 0

Possible failure on BAR 0. check dmesg from guest OS for corresponding hotplug BDF:

```
[10.874845] pci 0000:87:00.1: BAR 0: failed to assign [mem size
0x00100000]
```

> **Info**
>
> In this example, the hotplug PCIe BDF is 87:00.1. This value can be retrieved using "`lspci | grep -i virtio`" from the guest OS.

This can be normally resolved by adding "`pci=realloc`" in the Linux command line (grub).

BAR 14/15

Possible failure on other PCIe BAR. Check the dmesg from the guest OS for the corresponding hotplug BDF:

```
[ 2893.484281] pcieport 0000:10:01.0: bridge window [mem 0x00100000-
0x000fffff] to [bus 12] add_size 200000 add_align 100000
[ 2893.484285] pcieport 0000:10:01.0: BAR 14: no space for [mem size
0x00200000]
[ 2893.484287] pcieport 0000:10:01.0: BAR 14: failed to assign [mem
size 0x00200000]
[ 2893.484289] pcieport 0000:10:01.0: BAR 14: no space for [mem size
0x00200000]
[ 2893.484290] pcieport 0000:10:01.0: BAR 14: failed to assign [mem
size 0x00200000]
```

Info

In this example, the hotplug PCIe BDF is 10:01.0. This value can be retrieved using "`lspci | grep -i virtio`" from the guest OS.

- This is mostly due to there being insufficient BAR resources. Try to reduce the PF BAR size by performing the following from the BlueField side:

```
# mlxconfig -d 03:00.0 s PF_LOG_BAR_SIZE=0
```

- This can also be caused by the BIOS provider not reserving enough memory. Check the guest OS's dmesg for similar messages for the PCIe bus of the BlueField device:

```
[3.979061] pci_bus 0000:a0: root bus resource [mem 0x41c0800000-
0x41c10fffff window] (9M)
[3.979062] pci_bus 0000:a0: root bus resource [bus a0-bf]
[4.017770] pci 0000:a4:00.0:   bridge window [mem 0x41c0800000-
0x41c0ffffff 64bit pref] (8M)
```

```
[4.018243] pci 0000:a4:00.0: BAR 15: no space for [mem size
0x05800000 64bit pref] (88M)
[4.018245] pci 0000:a4:00.0: BAR 15: failed to assign [mem size
0x05800000 64bit pref]
```

- On the host, the prefetchable memory limit of the root bus (`a0`) is only 9 M. This means that all the devices under this bus (including BlueField) can only be allocated 9M prefetchable memory in total.
- The BAR 15 is the total prefetchable memory limit on the bridge (`a4`) of the device. The PCI bridge window of the BlueField for prefetchable memory is 8M, but the bridge requires 88M for its child device (BlueField). A fter several attempts, the PCIe bridge did not find sufficient IO memory to allocate for BlueField BARs. This can be solved by contacting the BIOS provider to provide enough memory to the PCI root.

Rescan

If the the hotplug operation from the BlueField Arm side is performed before the guest OS is up, and the virtio device is not found by the command " `lspci | grep -i virtio` ". Try to rescan from guest OS:

```
# echo 1>/sys/bus/pci/rescan
```

No Hotplug from BIOS

The server BIOS may not support hotplug device. This can be confirmed by looking at guest OS dmesg:

```
[8.209406] acpi PNP0A08:03: _OSC: platform does not support
[PCIeHotplug PME]
```

Try to enable hotplug from the BIOS:

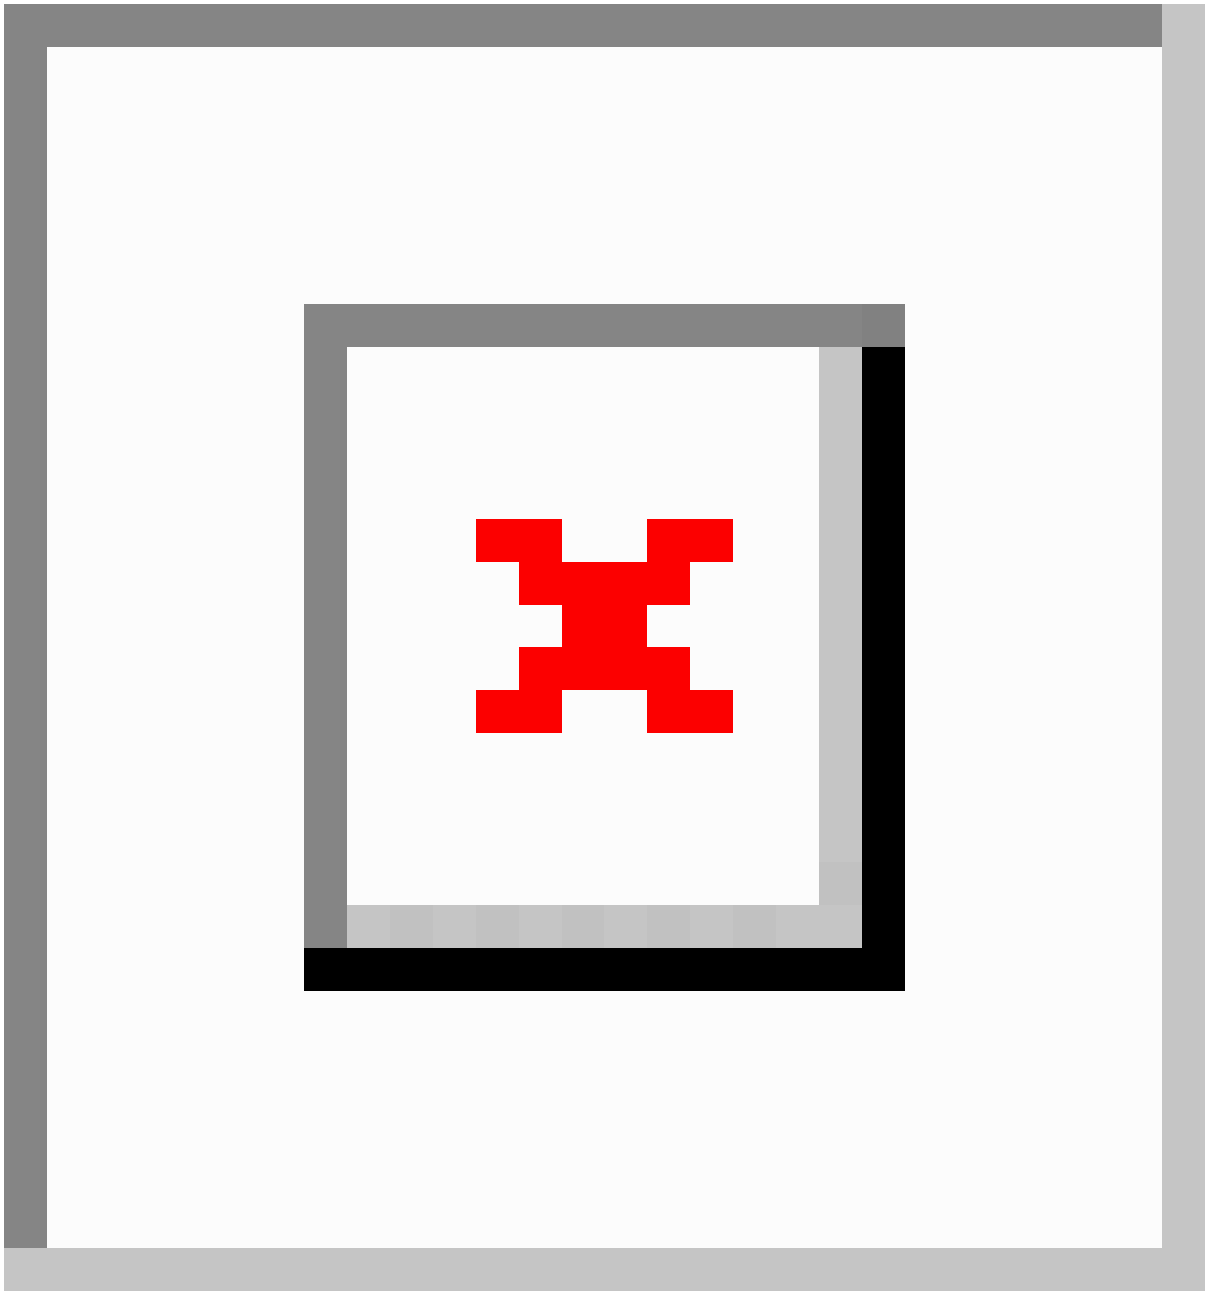

Force Hotplug

Guest OS may be running a kernel older than 4.19, the virtio device is not found by "`lspci | grep -i virtio`". Add the entry `pciehp.pciehp_force=1` to the grub command line.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF

ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2024, NVIDIA. PDF Generated on 11/12/2024